IN THE SPECIFICATION:

Please amend the carryover paragraph between pp. 7-8 of the Specification as shown below:

—The operation of the cache coherency mechanism is discussed below in conjunction with an example. Referring now to Fig. 4, the central switch 12 includes six interconnected switches 12_1, 12_2...12_6. A given switch 12_i connects to each of the other switches 12_j through switch ports 28. The switch also connects to eight associated group switches 20 through switch ports 18. Each group switch connects, in turn, to eight associated processors 22 through switch ports 48. The processor 22_3 is the current owner of the data block and the processors 22_1, 22_2, 22_4 and 22_7 hold copies of a data block 100 in their respective caches 26. The associated home node 30_1 which is connected to the central switch 12 ~~the associated~~ through home port 38 of switch 12_1, maintains the cache coherency directory 34. For ease of understanding, only certain network connections are depicted in the drawing.—

Please amend the paragraph starting at page 9, line 15 as shown below:

—When the home node 30 receives the message from the processor 22_4 (step 600), the home node enters the cache coherency directory 34 (step 602) and produces an

invalidate message that includes the arbitration masks 404_6, from the directory entry 300 for the data block 100 (step 604). The home node then sends the invalidate message over the home port 38 to the home switch 12_1 (step 606). The home switch decodes the first-level mask 406_1 and multicasts the message to the other switches 12_i through the ports 28 that correspond to the set bits 504_2, 504_4 and 504_6 of the mask. Thereafter, the switches 12_i decode the second-level arbitration mask 406_2 and direct the message through the ports 18 that correspond to the set bits 506_1, 506_3, 506_6 and 506_8. When the group of switches 20 receives the message, they consult their respective directories and/or routing tables to, as appropriate, direct the messages to the affected processors. Certain groups receive the message even though none of the affected processors are contained therein. These groups may simply ignore the messages. Alternatively, each group switch may locally broadcast the message, with the unaffected processors ignoring the message.—